

Measuring Productivity Dynamics with Endogenous Choice of Technology and Capacity Utilization: An Application to Automobile Assembly

Johannes Van Biesebroeck *

November 11, 2000

Abstract

During the 1980s, all Japanese automobile producers opened assembly plants in North America. Industry analysts and previous research claim that these transplants are more productive than incumbent plants and that they produce with a substantially different production process. We compare the two production processes by estimating a model that allows for heterogeneity in technology and productivity. We treat both types of heterogeneity as intrinsically unobservable. In the model, plants choose technology before production starts. They condition subsequent input decisions on this choice. Maximum likelihood estimation is used to estimate the unconditional distribution of the technology choice, output, and inputs. The model is applied to a sample of automobile assembly plants. We control for capacity utilization, unobserved productivity differences, and price effects. The results indicate that there exist two distinct technologies. In particular, the more recent technology uses labor less intensively and it has a higher elasticity of substitution between labor and capital. Hicks-neutral productivity growth is estimated to be lower, while capital-biased (labor-saving) productivity growth is estimated significantly higher, for the new technology. Using the estimation results, we decompose industry-wide productivity growth in plant-level changes and composition effects, for both technologies separately. Plant-level productivity growth is further decomposed to reveal the importance of capital-biased productivity growth, increase in capital-labor ratio, and returns to scale.

Key words: productivity growth, technological change, automobile industry, capacity utilization

JEL codes: C33, C34, D24, D92, L62.

*Department of Economics, Stanford University. Please send comments to jobv@stanford.edu. I wish to thank Frank Wolak for many helpful discussions and suggestions. I benefitted from comments from Lanier Benkard, Peter Reiss, and seminar participants at Stanford University. I am grateful to Mike Clune at the CCRDC for help with the data. Financial support from the Littlefield Fellowship, the GRO program at Stanford University, and the NSF under grant SES-9906403 is gratefully acknowledged. The research in this paper was conducted while the author was a Census Bureau research associate at the California Census Research Data Center. Research results and conclusions expressed are those of the author and do not necessarily indicate concurrence by the Bureau of the Census. This paper has been screened to insure that no confidential data are revealed.

Abstract

During the 1980s, all Japanese automobile producers opened assembly plants in North America. Industry analysts and previous research claim that these transplants are more productive than incumbent plants and that they produce with a substantially different production process. We compare the two production processes by estimating a model that allows for heterogeneity in technology and productivity. We treat both types of heterogeneity as intrinsically unobservable. In the model, plants choose technology before production starts. They condition subsequent input decisions on this choice. Maximum likelihood estimation is used to estimate the unconditional distribution of the technology choice, output, and inputs. The model is applied to a sample of automobile assembly plants. We control for capacity utilization, unobserved productivity differences, and price effects. The results indicate that there exist two distinct technologies. In particular, the more recent technology uses labor less intensively and it has a higher elasticity of substitution between labor and capital. Hicks-neutral productivity growth is estimated to be lower, while capital-biased (labor-saving) productivity growth is estimated significantly higher, for the new technology. Using the estimation results, we decompose industry-wide productivity growth in plant-level changes and composition effects, for both technologies separately. Plant-level productivity growth is further decomposed to reveal the importance of capital-biased productivity growth, increase in capital-labor ratio, and returns to scale.

Key words: productivity growth, technological change, automobile industry, capacity utilization

JEL codes: C33, C34, D24, D92, L62.

1 Introduction

By traditional measures, average productivity for the automobile industry has increased considerably since the 1980s. This increase coincided with the opening of Japanese assembly plants in North America. Because the number of vehicles produced per-worker in the Japanese transplants exceeds that in plants owned by American producers, researchers have concluded that industry-wide productivity increased because of the entry of more productive plants.

At the same time, it is claimed that the entrants produce with a different technology, dubbed lean, or modern, manufacturing. Lean production is associated with team work, less automation, flexible equipment, employee training, and increased flow in the production process.¹ The existence of two different technologies affects the evolution of productivity in the industry. Plants producing with different technologies will experience productivity growth at different rates and possibly different factor-biases. Some researchers also argued that imitation of the transplants by incumbent plants, has lead to additional productivity gains for the industry. The effect of these two different technologies on productivity measurement has not yet been addressed. Estimating production functions with random coefficients, or for limited samples of similar firms, has not yielded satisfactory results either (see Griliches and Mairesse (1990)).


In this paper, we investigate the rate and bias of total factor productivity growth for both technologies, taking the technology choice explicitly into account. Consistent with previous literature on automobiles production, heterogeneity in technology is explained by the existence of two distinct *systems* of production: lean and traditional. Plants choose between these two systems, but their technology choice is not directly observable to us.

The entry of Japanese transplants in the 1980s coincided with a large increase in vehicles produced per-worker. However, concluding from this that the entrants are more productive and that the labor productivity growth is driven by composition effects, is premature. Figure 1 plots the evolution of the vehicles per-worker statistic. After 1982 the statistic is calculated separately for plants already existing in 1982 and new plants. It suggests that the large increase is mostly situated with old plants. This could be due to an acceleration of productivity growth in the traditional technology, but it is equally possible that it represents old plants adopting the new technology. Accounting for input substitution complicates matters further. Differences in the factor-bias of productivity growth for the two technologies can also cause a different evolution of labor productivity. In order to make firm conclusions about the evolution of productivity in the industry and the underlying drivers, we have to disentangle the effects. In what follows, we will mostly concern ourselves with total factor productivity. Using the estimation results, we decompose the labor productivity evolution in the industry in different components using the primitives of our model.

Disentangling these effects is complicated by the basic unobservability of technology. Only for recent years, do we observe the exact production activities carried out at each plant. In addition, much of the variation between the two technologies is in the organization of work or the product flow through the plant, both of which are hard to measure. Furthermore, the probability that a

¹For example, many articles in the International Motor Vehicle Program (IMVP) have described in great detail how Japanese plants differ from their American and European competitors on several dimensions. The IMVP program has generated 328 working papers since 1986. The book by Womack et al. (1990) presents several findings of the program.

Figure 1: Evolution of vehicles produced per-worker for new and old plants



new plant chooses either technology is likely to vary across plants and over time. It is also likely that some plants, built before the new technology became available, switched technology at some later date. Therefore, we use an endogenous switching model of technology choice to investigate production decisions conditional on the technology choice.

Estimating a production function for the automobile industry is also complicated by several other factors that we control for. Large variations in capacity utilization distort the relationship between measured input levels and the actual services a plant derives from them. Price setting power makes deflated sales or value added an inappropriate measure of output. Unobserved heterogeneity in productivity levels leads to simultaneity bias, because plants choose inputs as well as output.

We estimate a separate production relation for lean and traditional producers, conditional on technology choice. To this end we collected data on a sample of automobile assembly plants in the United States. Input measures are obtained from the Longitudinal Research Data set (LRD), from the U.S. Bureau of the Census. It provides reliable input statistics because all plants are legally required to report the information. An advantage of the automobile industry is the existence of a well-defined unit of output: a vehicle. We collected information on actual production volumes to avoid using deflated sales or value added as output measure. This allows us to control for price effects and focus on actual production decisions. We also collected information to account for variations in input use. Capacity utilization in this industry varies considerably over time and among plants. We adjust for it by explicitly modeling the number of shifts a plant is operated. Finally, we obtained data on factors that influence the attractiveness of each technology. These include the type of vehicle assembled, ownership, and changeover dates.

Using a structural model of input choice and controlling for the choice of technology, we obtain consistent estimates for the parameters in the production function and the technology deci-

sion. We find the two estimated technologies to differ significantly. Ignoring them would bias the interpretation of the structural parameters, such as the rate of productivity growth. The traditional technology uses slightly more of both inputs and has constant returns to scale. The technology that we identify as lean production, experiences decreasing returns to scale, but increasing returns to shifts. The elasticity of substitution between capital and labor and the own-price elasticities are significantly larger for the lean technology. Both findings coincide with the notion that lean manufacturing is more flexible and relies less on standardization. The rate and bias of productivity growth for each technology are identified separately. Lean production is associated with a lower rate of Hicks-neutral productivity growth and a significantly higher rate of capital-biased, labor-saving, productivity growth. Labor productivity growth for traditional producers is mainly driven by increases in total factor productivity growth. For lean producers, capital-biased productivity growth and an increase in capital per worker are the two most important contributors to labor productivity growth.

We also find, not too surprisingly, that switching to lean technology has become more likely over the years. This is potentially driven by the proliferation of models, produced in small production runs. A technology switch is also more likely in changeover years, when switching costs are lowered. Some observers speculated that entry of a new technology would lead to an increase in productivity growth for mass producers, making them catch up with lean production. The results indicate that many mass producers have adopted lean production, instead of improving the mass production technology. Decomposing industry-wide labor productivity growth illustrates the importance of plant-level labor productivity growth by lean producers. It also shows that relocation of resources from traditional to lean producers accounts for a significant part of the observed growth rate.

The next section provides an overview of how we consistently estimate both production technologies. It also describes the timing of decisions and productivity shocks in the model. Section 3 contains a description of the data and motivates the existence and effect of two technologies. The four different stages of the model are described in more detail in section 4. Section 5 derives the likelihood function and contains the estimation results. In section 6 we derive some conclusions about the evolution of productivity growth in the industry.

2 Features of the model

2.1 Empirical strategy

A brief digression on the interpretation of a production function reveals many of the estimation difficulties. The technology a plant produces with —as captured by the production function— is a relationship that determines the maximum output that can be obtained from a bundle of physical inputs. Some plants will produce more output with the same amount of inputs. These plants can still be considered to be producing with the same technology, if we allow for a different level of productivity. Input substitution possibilities are identical, but the production frontier is shifted outward radially. In addition, technology can shift over time. A radial shift that affects all inputs identically is called Hicks-neutral productivity growth. In this industry, it is often assumed that a significant part of technological progress comes through improved machinery and equipment: capital-biased (or labor-saving) productivity growth. In order to measure the shape and shift of the

production function we want to observe physical amounts for inputs and output and the individual productivity level for a sample of plants that we know to produce with the same technology. Each of these four elements —technology, productivity, output, and inputs— pose specific problems.

In order to compare productivity growth for the lean and traditional technology we need to obtain consistent estimates of the parameters in each production function.² Simply regressing output on inputs is unlikely to produce this, because we do not observe the technology choice directly. At one extreme we could postulate that all plants in our sample produce with the same technology. All observed heterogeneity would then be attributed to measurement or sampling error. At the other extreme we could assume a different technology for each plant. This assumption would render estimation nearly infeasible. In addition, Diamond et al. (1978) show that it is impossible to identify the bias in productivity growth from the elasticity of substitution using only time-series variation. We take an intermediate position by allowing two —but only two— technologies. Using observable characteristics that determine the technology choice we predict the probability that a plant produces with each technology.

The second piece of information that is intrinsically unobservable is the productivity level of individual plants. Given that inputs as well as output are chosen by the plant we have a source of simultaneity bias, making least squares estimation of the production relation inconsistent. Instrumental variables are the traditional solution. Blundell and Bond (1998) demonstrate that in the context of the production model, it is very hard to come up with powerful instruments. One solution is to use a behavioral equation to obtain an expression for the unobserved productivity in terms of observable variables. For example, Olley and Pakes (1995) invert the investment function to substitute the productivity term from the production relation.

In order to identify both the technology choice process and the production functions, we rely on distributional assumptions for the unobservable variables. These assumptions are critical for the construction of the model and also dictate a maximum likelihood approach. We predict the probability for each technology for each plant-year. The distribution of inputs and output will depend on the technology a plant has chosen. In the maximum likelihood framework, we put both together, to obtain the unconditional distribution of the endogenous variables.

Where possible, we collected data to measure physical output and inputs directly. Using data on output quantities, shifts, and hours worked, we control for price effects and variations in capacity utilization. Where physical inputs are unavailable we use stylized facts about the industry to model input choice.

2.2 Timing


The equations in the model are derived in section 4, but we give an outline of the model here. We will specify a production function, dependent on capital, labor, materials and shifts. We observe plants choosing the exact number of shifts to operate a plant during the year. From the production relation and optimal input choices, we derive three estimating equations, conditional on technology. Two different sets of coefficients are estimated, each representing one technology. A

²The level of productivity is only defined relative to a particular production function. Therefore we can only compare the productivity level of plants that face the same technology. Productivity growth, on the other hand, is always well-defined. It measures the shift of a production frontier over time. Productivity growth will be the focus of our analysis.

model of technology choice generates a predicted probability for each technology, which is used to construct the likelihood function. Unobserved heterogeneity in productivity has two components, one constant over time and one variable.

Figure 2 shows the time line of decisions and indicates when the errors are realized. The first line describes decisions taken when the plant is built, before the current production year. After some preliminary choices, that we do not consider explicitly, the plant makes a technology choice (i), traditional or lean. The constant component of productivity (ω_j) is realized before any production decision is made. It is known to the plant, but we do not observe it.

Figure 2: Timing of decisions and errors



The second line describes the time line of events in each year. First, the plant has the option of switching technology from traditional to lean. This decision is modeled using a reduced-form approach. Then output (Q) is decided. We assume that a plant gets the production requirement handed down from headquarters. It minimizes costs subject to the production constraint. The second decision a plant makes is investment (I). Investment (and capital, (K)) is potentially a function of the individual productivity level but not of any other random variable in the model. Before the start of production, the plant learns the exact productivity of the capital stock (ϵ^k), which determines the efficiency units of capital (\tilde{K}) used in production. The optimal choice of shifts (S) and labor (L), the third decision, generates the first estimating equation. At the same time it provides us with an expression for capital productivity (ϵ^k) in terms of observable variables.

The fourth and final decision concerns actual production. Actual output (Q) will differ from planned output (\bar{Q}), because of an ex-post shock to production (ϵ^q), realized after labor and shifts are chosen. Finally, material input is proportional to the actual output produced. We do not observe the volume of materials (M), only the value of materials (\hat{M}). These are related by an index that represents quality upgrading in components and includes a stochastic term (ϵ^m). Production generates two additional estimating equations. The second equation we estimate relates output to capital, labor, and shifts. The third equation is the relationship between material input and output.

There are two sources of unobserved productivity differences in the production function. We control for both to avoid simultaneity bias. The first component (ω_j) is assumed to be constant over time. It is captured by plant-fixed effects. The second component (ϵ^k) is variable and represents

a shock to capital productivity. Much of the productivity improvement in this industry is caused by technological advances in machinery. We assume that a plant cannot predict this perfectly when deciding investment. It chooses variable inputs after observing the actual productivity of the capital stock. Using the optimality condition for the choice of variable inputs, we can express ϵ^k in terms of observable variables. The only error term in the production function we do not observe is a shock to production (ϵ^q) that makes the plant miss its planned output. We assume that this random shock is realized after inputs are chosen and independent of both productivity components and input levels.

The following table summarizes the primitives of the empirical model ($i = L, T$).

<u>decision</u>	<u>equation</u>	<u>errors</u>
(1) i	$\Pr(i = T)$	\rightarrow reduced-form for technology choice
(2) K	$= g(\omega_j, \dots)$	\rightarrow (not estimated)
(3) $g_i(\frac{L}{S})$	$= g_l(w, K, t; \beta_i) + \epsilon_i^k$	$\rightarrow \epsilon_i^k = g_i(\frac{L}{S}) - g_l(\dots)$
(4a) Q	$= g_q(S, L, K, t, \epsilon_i^k, \omega_j; \beta_i) + \epsilon_i^q$	$\rightarrow \epsilon_i^k$ from (3), ω_j is plant-fixed effect ϵ_i^q independent of inputs, ϵ_i^k , and ω_j
(4b) \hat{M}	$= g_m(Q, t; \beta_i) + \epsilon_i^m$	$\rightarrow \epsilon_i^m$ realized at the very end

The technology decision (1) is used to derive the probability a plant produces with the traditional technology in each year. The labor-shift decision (3), production (4a), and materials (4b) provide the three estimating equations. The equations imply a distribution for each endogenous variable conditional on technology. The likelihood function for the unconditional joint distribution is obtained by multiplying the probability for each technology and the conditional distributions for the three endogenous variables. We assume that the three error terms — ϵ_i^k , ϵ_i^q , ϵ_i^m — are independently distributed. In section 4 we describe the model in greater detail and we derive the equations explicitly. First, we introduce the data and motivate the existence and effect of two technologies in the automobile industry.

3 Industry Characteristics

3.1 Data

From the preceding discussion it is clear that the data needed is threefold: variables characterizing the technology choice, output quantities, and input levels, adjusted for intensity of use.

The principal data source is the Longitudinal Research Data set (LRD) constructed by the Center for Economic Studies at the U.S. Bureau of the Census. The data is taken from plant responses to the Annual Survey of Manufactures and the Census of Manufactures. Observations are plant-years and plants are linked over time. Coverage includes all plants with SIC code 3711 (motor vehicles and car bodies) for their main products. The data spans the years 1963, 1967, and 1972-96.

Industry publications are used to supplement the LRD. These cover a smaller number of companies and plants. Only statistics for plants owned and operated by one of the large automobile

or truck companies are available. Omitted plants are owned by smaller firms and specialize in converting cars to limousines, trucks to campers or they only make car or truck bodies. In addition, the LRD contains some engine or component plants that produce a large number of bodies or completed vehicles only sporadically.³ 74% of the observations in the LRD with SIC code 3711 could be matched to the data from other sources. These plants represent 94% of the employment in the industry.

This data set provides reliable and complete input statistics. The labor input measure we use is defined as total hours worked at the plant. Hours worked by non-production workers are imputed using their relative wage. Unionization and volatility in production load make that companies have a number of temporarily unemployed workers on their payroll. Only actual hours worked are counted in labor input. The amount paid to temporarily unemployed workers is included in the labor costs. Because we observe whether plants are unionized or not, we can account for this in the empirical model. Capital input is constructed from book values. An alternative measure, using the perpetual inventory method, yielded almost identical results. It is deflated using the capital goods deflator for the industry from the NBER productivity data set. Material input includes raw materials and intermediate products, fuels, and electricity. All are scaled by the appropriate deflator from the NBER data set. In principal, we could include energy separately in the production function. This was not done because it is smaller than 1% of costs for almost all plants and included in raw materials for some. Table 1 contains summary statistics for the relevant variables.

The second piece of information needed is output. Most productivity studies use deflated sales or value added as output measure, because actual production volumes are not generally available. If a firm has price setting power, price changes will erroneously be interpreted as productivity changes. For example, if a firm produces subject to an inelastic demand, it can increase sales by raising the price. Deflation by an industry-wide price index does not capture these individual price movements. Because output and inputs do not change—or even decrease—this increase in “output” will be interpreted as a productivity gain. In a concentrated industry like automobile assembly, price setting is likely to be important. One example is the sale of identical vehicles by different firms. All American producers have joint-ventures with Japanese partners, assembling vehicles jointly. Two identical products are sold under different nameplates. The Japanese model invariably fetches a higher price.⁴ We collected information on the number of cars and light trucks produced by each plant. For the years 1985-1996 this data was obtained directly from Ward’s Automotive. For the preceding years two data sources in Ward’s Automotive Yearbook are matched.⁵

The input and output statistics are augmented by variables that allow us to distinguish both technologies. These include dummies for the type of vehicle produced and a dummy for Japanese

³Plants are classified according to the industry category of their main product. Some engine or component plants assemble a limited amount of vehicles as well, which in some years can make up a large part of sales. This can result in these plants being classified in SIC industry 3711 in some years.

⁴This price difference can be substantial. NUMMI, the joint-venture between GM and Toyota in Fremont, CA, produces the Chevrolet Prizm and Toyota Corolla. Both models are identical, assembled from the same components on the same assembly line. On average, the Prizm is sold for \$3000 less than the Corolla, which is about 20% of the average retail price. Some of this represents lower profits for Chevrolet dealers, but most of it is a factory rebate. Using deflated sales or value added as output measure makes NUMMI look much more productive assembling Corollas.

⁵Details about the calculations are available upon request. As a robustness check the calculated production for 1985 is compared with the information obtained directly. The correlation was a reassuring 0.99.

Table 1: Summary statistics for the automobile assembly industry (1963,1967,1972-96)

variable	mean	standard deviation
<u>Inputs:</u>		
total hours worked per-shift	22386	17562
total employment	4332	3133
production workers (% of total)	0.86	0.05
book value of capital (% of sales)	.334	.293
materials-sales ratio	0.75	0.10
energy-sales ratio	0.006	0.004
<u>Output:</u>		
cars produced	181780	92076
light trucks produced	130547	81637
total vehicles produced	196902	100177
<u>Other variables:</u>		
only-cars dummy	.57	.49
only-trucks dummy	.23	.42
cars and trucks dummy	.20	.40
Japanese ownership dummy	.05	.22
changeover dummy	.04	.18
annual capacity	237095	83904
shifts operated (per-year)	424	114
union dummy	.97	.18
number of observations	1358	
number of plants	78	
number of firms	17	

Sources: Longitudinal Research Data set, Bureau of the Census, 2000; Ward's Automotive; and Automotive News weekly magazine (various years).

ownership. A dummy for changeover years will be used as a proxy for switching costs.

Information on shifts is collected to account for capacity utilization. A major concern for productivity measurement in this industry is the volatility in capacity utilization. The Harbour report (1999) calculates assembly plant utilization rates for 1998, a record production year, between 34% and 148%.⁶ Even though most of the capital cost is sunk after it is installed, many plants choose to remain idle for part of the year. The tradeoff a plant faces is to run few shifts, with many workers on each, or run a lot of them, with fewer workers. The number of shifts to operate the plant is an explicit choice the plant makes, resulting in endogenous variation in capacity utilization. From Automotive News weekly magazine we obtained the number of weeks a plant ran overtime, worked on Saturdays, and the number of weeks a plant closed for vacation, for inventory adjustment, and for retooling and model changeover. Using these five measures the total number of shifts the plants operated each year can be computed.


⁶Capacity utilization can exceed 100% since plants can run more than 2 shifts a day, 5 days a week, while capacity is calculated for regular operation.

3.2 Existence and effect of two technologies

As mentioned earlier, we do not observe the technology choice directly and assume there are two types. The trade press takes this stance by drawing a sharp distinction between lean and mass production. Milgrom and Roberts (1992) provide theoretical support for limiting the technologies to only two types. They describe modern manufacturing as a set of activities that exhibit complementarities. The marginal product of adopting the new technology for one activity is increasing in adoption on other dimensions. This makes intermediary systems that are composed of elements from the traditional and modern systems unstable.

Figure 3 provides additional evidence for the existence of two technologies and the possibility that plants switch between the two. The left panels plot non-parametric densities for the capital-labor ratio for the first five years of the sample.⁷ The right panels plot the same graphs for the last five years of the sample. The top panels contain the ratio for all plants and the bottom graphs are limited to plants that remained in the sample for the entire sample period, from 1963 to 1996. Comparing left and right panels, we see that the ratio has a bimodal distribution in both time

Figure 3: Non-parametric distribution for the average capital-labor ratio for different groups of plants




periods. In the early years, most plants choose the low capital-labor technology, the left mode. In later years, most plants prefer the technology with a higher capital-labor ratio, leading to an increase in the right mode. In particular, this is true for the bottom panels, containing the same plants in both periods. We interpret this as evidence of plants switching technology.

⁷These graphs can be interpreted as smoothed histograms. An observation is the average capital-labor ratio for a plant over the relevant five year period. Confidentiality considerations preclude us from reporting the underlying statistics directly.

We illustrate with an example that failure to control for the existence of different technologies will bias productivity growth measures. Figure 4 plots the unit isoquants for two types of technologies in input space. The producers with the technology that increased in popularity, the new technology, have a steeper isoquant (the solid line), choosing a higher capital-labor ratio at the same factor prices. Producers with the older technology face a different tradeoff between labor and capital, captured by the dashed line. If both types of plants are pooled, the estimated technology will lie between the two existing ones and have the shape of the dotted line. Take a plant with an initial production plan at P_0 and a production plan P_1 in some later year. Without knowledge of the input tradeoff the technology allows, it is impossible to know the growth in total factor productivity the plant experienced. We have to separate movement along the isoquant from the shift in the function. If we estimate only one production function for the pooled sample, productivity growth will be underestimated for plants with the traditional technology. Actual productivity growth is $2P_1/02$, although it is estimated to be $1P_1/01$. For producers with the new technology, productivity growth is overestimated. It is actually zero, although it is estimated to be $1P_1/01$ as well. If we believe heterogeneity in technology to exist, we have to control for it to measure productivity growth correctly.

Figure 4: The existence of two technologies will lead to biased measures of productivity growth.



4 A Model of the Decision Process

Now we analyze the four decisions —technology choice, investment, variable inputs, and production— in greater detail. Earlier choices take optimal decisions in later stages into account. Later

choices are conditional on the outcomes of earlier decisions. We start with the last decision, production, and work our way backwards. The production decision will yield two equations we estimate, one for materials and one for output. Secondly, the choice of variable inputs will yield the third equation. It also generates an expression for the capital productivity shock, needed to estimate the production function consistently. Thirdly, we discuss the investment decision. Nothing is estimated from this, but we derive three interesting properties. It also indicates that capital will be a function of the constant component of productivity, but not of any other unobservable in our model. The first three decisions are conditional on the fourth and final decision, technology choice. This is modeled in a reduced-form way and used to construct the likelihood function.

4.1 Fourth decision: Production

Ideally we would like to observe the physical input of material in production. Knowledge about the assembly activities carried out at each plant could be taken into account in the production function, making sure it is a stable relationship. We only observe the value of material input, which displays two important trends. The first observation is the large increase in (the value of) materials per-vehicle over the sample period ($\frac{P_m M}{Q}$). At the same time the material-sales ratio remained virtually constant ($\frac{P_m M}{P_i Q}$).

Quality upgrading of vehicles, through higher quality components, provides one explanation consistent with both observations. More recent cars use better fabric, better quality paint, more powerful brakes, etc. The amount of intermediary inputs per-vehicle remained constant, but the quality and (real) price increased over time. This is consistent with the observed increase in price for the final product. The real price per-vehicle (in 1987 \$) increasing from \$6000 to \$13000 over our sample period.

An evolution in outsourcing activities provides an alternative explanation for the two stylized facts. The use of more material inputs in the assembly process can explain the upward trend in the material-vehicle ratio. In order to reconcile this with the constant material-sales ratio, the price increase for vehicles has to outpace the price increase for materials. We do not find evidence for such a trend at all. This explanation faces an additional problem if the capital-labor combination to assemble different sub-components, of equal value, differs. In that case the production function depends on exactly which activities are outsourced and no stable material aggregate exists. Without observing the actual material inputs, it is impossible to remedy this problem.

We adopt the first interpretation, quality upgrading with a constant amount of material inputs per-vehicle. The trade press describes the trend to increased outsourcing as something operating at the firm level. Instead of receiving components from in-house suppliers, more is bought from outside firms. Both options have little impact on the activities carried out at the assembly plant. If we compare the volatility in the material-sales ratio with the volatility for capital or labor input, it is almost negligible, providing additional support for the assumption of constant material input.⁸

⁸We use the following formulas to evaluate the volatility of the material-sales ratio over time: $\frac{1}{J} \sum_{j=1}^J \frac{\frac{1}{T} \sum_t (\frac{\dot{M}}{S} - \frac{1}{T} \sum_t \frac{\dot{M}}{S})^2}{\frac{1}{T} \sum_t \frac{\dot{M}}{S}}$ and across plants: $\frac{1}{T} \sum_{t=1}^T \frac{\frac{1}{J} \sum_j (\frac{\dot{M}}{S} - \frac{1}{J} \sum_j \frac{\dot{M}}{S})^2}{\frac{1}{J} \sum_j \frac{\dot{M}}{S}}$, and similarly for labor and capital. Using these measures we see that capital-sales is 6 times more volatile than material-sales. Labor-sales is 4 times more volatile. In addition, the difference between material volatility across plants is hardly higher than the volatility

We model the production function as Leontief, allowing no substitution between materials, on one hand, and capital and labor, on the other hand. The functional form is motivated by industry practice. For most components, there is hardly any scope for substitution. Engines, transmissions, and electrical components are outsourced by all plants. Most other components, such as shock dampeners, seats, wheels, etc., are also outsourced and simply installed at the assembly plant. The fixed coefficient for material input is allowed to vary between technologies. Production for plant j producing at time t with technology i is governed by the following relationship:

$$Q_{jt} = \min [\alpha'_{i0} M_{jt}^{\alpha_{im}}, \bar{Q}_{jt} e^{\epsilon_{ijt}^q}] \quad (5)$$

The first part of the function relates the output quantity to the volume of material inputs. Because we do not observe material input directly, we link the value of materials to output using a price index. The index captures price and quality upgrading of components and includes a stochastic component. For example, automobile producers have the choice between more expensive, but higher quality, disk brakes and cheaper drum brakes. During our sample period, manufacturers substituted most drum brakes for disk brakes. The component cost per vehicle increased and it is not captured by price deflation because both brakes are different goods. At the same time, the amount of material input per-vehicle remained constant (four brakes per-vehicle) and assembly time also did not change.

The price index (P_i) links the observed value of materials (\hat{M}), measured in 1987 dollars, to the unobserved amount of materials used (M):

$$\begin{aligned} \hat{M}_{jt} &= P_{ijt} M_{jt} \\ P_{ijt} &= e^{p_{i0} + \theta_{im} t + \epsilon_{ijt}^m}. \end{aligned}$$

θ_{im} is the average price increase for material inputs. ϵ_i^m is a stochastic component that is unobservable to the econometrician.⁹ The amount of materials is not observed, but we can obtain it from the production function. If all inputs have a positive price, both parts in the Leontief function will hold with equality. Taking logarithms, rearranging, and substituting M gives the first equation we estimate,

$$(4b) \quad \hat{m}_{jt} = \alpha_{i0} + \frac{1}{\alpha_{im}} q_{jt} + \theta_{im} t + \epsilon_{ijt}^m.$$

Small cap variables indicate the logarithm of a variable.

The second part in the production function indicates that measured output (Q) will differ from planned output (\bar{Q}), because of the realization of a shock to production (ϵ^q). Planned output is a function of capital, labor, shifts, and productivity. The fixed-coefficient technology for materials makes it reasonable to assume that materials are determined after the realization of the production shock (ϵ^q).¹⁰ We define the per-shift production function as

$$\bar{Q}_{jt} = S_{jt}^{\alpha_{is}} f_i\left(\frac{L_{jt}}{S_{jt}}, K e^{\theta_{ik} t + \epsilon_{ijt}^k}\right) e^{\theta_{in} t + \omega_j}. \quad (6)$$

over time for a given plant.

⁹It does not matter whether a plant observes ϵ_i^m or not, because it has no control over the amount of materials to use. Output is determined exogenously (at the firm level) and material input is linked to output through a constant-coefficients technology.

¹⁰It is straightforward to adjust the model making the choice of material input precede the production shock. The only change to the estimation would be the introduction of a positive correlation between the errors in the production and material equation.

This incorporates Hicks-neutral productivity differences among plants (ω_j) and neutral productivity growth (θ_{in}). Capital-biased productivity differences are captured by a plant-specific shock (ϵ_i^k) and a shift over time (θ_{ik}). A plant observes all four factors.

Output per-shift is a function of labor per-shift and efficiency units of capital. The shape of $f_i(\cdot)$ determines the technological substitution possibility between capital and labor for technology i . To obtain total production, we multiply with a scale factor for the amount of shifts a plant is operated ($S_{jt}^{\alpha_{is}}$). The α_{is} coefficient can be smaller than one. For example, running more shifts can reduce maintenance time, leading to more machine breakdown and lower production per-shift. It can be larger than one if there are positive spillover effects between shifts. Examples are shared overhead or reduction in start-up time.

Substituting the expression for planned output in the production function gives the second equation we estimate,

$$(4a) \quad q_{jt} = \alpha_{is} s_{jt} + \log f_i\left(\frac{L_{jt}}{S_{jt}}, \hat{K}_{ijt}\right) + \theta_{int} + \omega_j + \epsilon_{ijt}^q,$$

where \hat{K}_i are efficiency units of capital ($\hat{K}_i = K e^{\theta_{ik}t + \epsilon_i^k}$). The equations (4b) and (4a) describe the distribution of materials and output conditional on technology choice. All coefficients and the error term are technology-specific. The only element unobserved to plants is the ex-post shock to output. We control for ω_j , using plant dummies, and for ϵ_i^k , using an expression derived from optimal input choices (in the next section). We use the translog specification for the labor-capital aggregate,

$$\log f_i = \alpha_i + \alpha_{il}l + \alpha_{ik}\hat{k}_i + \frac{1}{2}\beta_{il}l^2 + \frac{1}{2}\beta_{ik}\hat{k}_i^2 + \beta_{ilk}l\hat{k}_i + \theta_{int} + \omega_j + \epsilon_i^q,$$

where l is the logarithm of labor per-shift.

4.2 Third decision: Choice of variable inputs

Most productivity studies using the translog production or cost function use Shepard's lemma to derive the factor-share equations, to aid in identification. We have the additional need to obtain an expression for the shock to capital productivity in terms of observable variables. The introduction of shifts as choice variable and the fact that plants choose labor and capital at a different time, makes it necessary to solve the cost minimization explicitly.

Every period a production plan is handed down from headquarters. Capital, output and technology are not decision variables at this point. The plant chooses labor and shifts to satisfy the production requirement and to minimize variable costs, which are two-fold. First, wages are proportional to the number of shifts and the average number of hours worked on a shift. Second, there are fixed and variable costs associated with operating a shift. The existence of labor unions complicates the tradeoff between labor and shifts the plant makes. Unionized plants save less when they reduce labor input. Labor contracts negotiated in this industry specify that a percentage of the normal wage is paid even when plants are idled. Because plants only take the variable portion of the wage into account, we multiply the observed wage by a factor that lies between zero and one (δ). For non-unionized plants, this fraction is normalized to be one.

The costs associated with operating a shift are not observed directly. We estimate them with a fixed component (μ_i) and a component proportional to capital (ρ_i).¹¹ We observe many

¹¹There is no unique solution if returns to scale, for the per-shift production function, are equal to α_{is} . With excess

plants idling the capital stock for part of the year. This indicates a positive marginal cost of operating the capital stock. Reasons for this can be complementary inputs (labor and energy), depreciation in-use, or maintenance cost. Both components are equal across plants, but possibly different for each technology.

A plant solves the following cost minimization problem:

$$\begin{aligned}
\min_{\{L, S\}} \quad & S \times (w \delta \frac{L}{S} + \rho_i K + \mu_i) \\
\text{s.t.} \quad & Q = S^{\alpha_{is}} f_i(\frac{L}{S}, K e^{\theta_{ik} t + \epsilon_i^k}) e^{\theta_{in} t + \omega_j + \epsilon_i^q} \\
& E(Q) \geq \bar{Q} \\
& \delta = 1 \quad \text{if plant is not unionized} \\
& \delta \in [0, 1] \quad \text{if plant is unionized} \\
& K \leq \bar{K}.
\end{aligned} \tag{7}$$

In the objective function, w is the observed wage rate and δ is the fraction of the wage that is not paid when a plant is idled. $\frac{L}{S}$ is the average hours worked per-shift by all employees. K , the capital stock, is fixed at this point. S is the number of shifts the plant is operated over the entire period. Only the relevant part of the production function is repeated.

This model allows for the large under-utilization of capital often noted in the automobile assembly industry. If capital has a positive operating cost ($\rho_i > 0$), it can be advantageous to idle the plant some shifts. This is especially relevant if output requirements change after the capital stock is fixed, but before variable inputs are chosen. It also applies if a plant chooses variable inputs more frequently than investment. A manufacturer can employ more workers on each shift and run fewer of them to save on capital depreciation, maintenance, and energy. This can only be accomplished by running the assembly line at a higher jobs-per-minute (JPM) rate, made possible by the increased labor input. Since reported capacity numbers for the industry are calculated as potential output, assuming ten shifts per-week and the initially reported JPM rate, this substitution behavior will show up as lowered capacity utilization. The tradeoff between labor and capital is determined by the operating cost per-shift (ρ_i), the returns to shifts (α_{is}), and the elasticity of substitution between capital and labor. All three factors are identified separately in our model.

To obtain an equation to estimate, we solve the minimization problem and rework the first-order conditions. The Lagrangian we minimize is

$$\mathcal{L}(L, S) = w \delta L + (\rho_i \bar{K} + \mu_i) S + \lambda [\bar{Q} - S^{\alpha_{is}} f_i(\frac{L}{S}, K e^{\theta_{ik} t + \epsilon_i^k}) e^{\theta_{in} t + \omega_j}],$$

where L and S are the decision variables. The fixed cost per-shift and scarcity of resources will guarantee that $K = \bar{K}$ and $E(Q) = \bar{Q}$.¹² The first-order conditions are

$$w \delta = \lambda e^{\theta_{in} t + \omega_j} S^{\alpha_{is}-1} f'_{iL/S} \tag{8}$$

$$\rho_i \bar{K} + \mu_i = \lambda e^{\theta_{in} t + \omega_j} S^{\alpha_{is}-1} (\alpha_{is} f_i - f'_{iL/S} \frac{L}{S}). \tag{9}$$

capacity, a plant can produce the same output in t shifts, using L hours per shift, or using all tL hours in one shift and employing t times as much capital. Introducing a fixed cost per-shift guarantees that a plant will always use the entire capital stock if it operates a shift.

¹²The only variable in the problem a plant does not observe at this stage is the ex-post shock to production, ϵ_i^q .

Dividing the two equations and rearranging gives

$$\underbrace{\frac{w \delta L}{w \delta L + S(\rho_i \bar{K} + \mu_i)}}_{\lambda_i} = \frac{1}{\alpha_{is}} \frac{\partial \log f_i(\frac{L}{S}, \hat{K})}{\partial \log(\frac{L}{S})}. \quad (10)$$

The expression on the left is the labor share of variable cost. The variable cost share is not directly observable and we have to estimate the parameters ρ_i , μ_i , and δ . Equation (10) indicates that the optimal labor-share ratio does not depend on the plant-specific productivity or the shock to production.

Using the translog specification for f_i we can write

$$(3) \quad \frac{w \delta \frac{L}{S}}{w \delta \frac{L}{S} + \rho_i \bar{K} + \mu_i} - \beta_{il} \log \frac{L}{S} = \alpha_{il} + \beta_{ilk}(k + \theta_{ik}t + \epsilon_i^k).$$

The dependent variable, $\frac{L}{S}$, is a nonlinear function of the disturbance, ϵ_i^k . We use this identity to substitute for ϵ_i^k in the production function, which allows for consistent estimation of equation (4a). In addition, we estimate equation (3) directly, to aid in identification.

4.3 Second decision: Investment

We briefly introduce the assumptions made to estimate the production function and derive some interesting theoretical results. A plant decides investment taking output and technology as exogenous and anticipating optimal choices of labor and shifts. If investment can be negative and capital-biased technological change is disembodied, it is not necessary to consider later years.¹³ The manufacturer faces the following optimization problem

$$\begin{aligned} \min_{\{I\}} \quad & rI + E[G(w, K, \bar{Q})] \\ \text{s.t.} \quad & Q = S^{\alpha_{is}} f_i(\frac{L}{S}, K e^{\theta_{ik}t + \epsilon_i^k}) e^{\theta_{in}t + \omega_j + \epsilon_i^q} \\ & E(Q) \geq \bar{Q} \\ & G = w\delta L^* + (\rho_i K + \mu_i)S^* \\ & K_t = (1 - d)K_{t-1} + I_t, \end{aligned} \quad (11)$$

where r is the user cost of capital and G is the variable cost function. L^* and S^* are the optimal values for labor and shifts derived from the variable input optimization. Both are functions of capital as well.

The Lagrangian to minimize is

$$\begin{aligned} \mathcal{L}(I) = \quad & rI + E[G(w, (1 - d)K_{t-1} + I, \bar{Q})] \\ & + \lambda' \{ \bar{Q} - S^{\alpha_{is}} f_i(\frac{L}{S}, K e^{\theta_{ik}t + \epsilon_i^k}) [(1 - d)K_{t-1} + I] e^{\theta_{in}t + \omega_j + \epsilon_i^q} \}, \end{aligned}$$

which gives the first-order condition for investment

$$r + E\left[\frac{\partial G^*(K)}{\partial K}\right] = p^e E[S^{\alpha_{is}} f'_{iK}] e^{\theta_{in}t + \omega_j}. \quad (12)$$

¹³Disinvestment is not a rare occurrence in this industry. Nearly 55% of the observations in the sample sell off used equipment or buildings. Gross investment is negative in 3.5% of the cases.

The first term on the left-hand side is the usual user cost of capital. This captures time-depreciation, change in valuation of the asset and interest cost. The second term is the marginal in-use cost for capital. It includes the cost of operating shifts and takes the change in optimal labor and shifts into account. On the right-hand side we use two simplifications. The shadow cost for the production equation is the value of one unit of output. We substitute the expected price for the Lagrange multiplier because the output level, and hence the price, is exogenous to the plant. We also substituted $E(\epsilon_i^q) = 0$. This error term is independent of all other errors and the decisions of the plant.

Because labor, shifts, and the marginal product of capital depend on ϵ_i^k in a nonlinear way, the first-order condition does not simplify further. A plant will choose investment in order to satisfy

$$(2) \quad \frac{E[(r + \frac{\partial G^*(K)}{\partial K})K]}{p^e \bar{Q}} = E\left[\frac{\partial \log f_i}{\partial \log K}\right].$$

On the left is the expected total cost of capital as a percentage of planned sales over the entire year. On the right is the logarithm of the expected marginal product of capital. Compared to the usual capital share equation, capital cost is now the sum of a *variable* cost, which depends on capacity utilization, and a *fixed* (user) cost, which is sunk after capital is installed.

Estimation of equation (2) would be possible if f_i takes the Cobb-Douglas form. In that case the marginal product of capital is a constant and it greatly simplifies the expected variable cost. A major disadvantage would be that it fixes the elasticity of substitution between capital and labor to unity. It would also make it impossible to identify the factor-bias in productivity growth. Using the translog function instead, we do not estimate the capital demand and forgo the ability to separate the fixed and variable components of capital cost (r and ρ). The appendix derives three interesting properties of the investment decision. K is a function of ω_j and we account for it by using plant-fixed effects in the estimation of the production relation. It also depends on the parameters of the distribution of ϵ_i^k , but not on the actual realizations.

4.4 First decision: Choice of technology

We do not observe the outcome of the technology choice directly. In order to estimate the parameters in the production function, we need to control for it. The previously derived input decisions are all conditional on the technology choice. If we can predict the probability that a plant is producing with either technology, we obtain the unconditional distribution for the endogenous variables.

We have to make specific assumptions to determine the probability a plant is producing with each technology. Beard et al. (1991) take an agnostic position, using an exogenous switching model. The error term in their cost function is distributed with a mixture of normal densities. They estimate the coefficients in two cost functions and a parameter capturing the incidence of each technology in the sample. The bimodal pattern in Figure 3 is consistent with the existence of two technologies. In addition, the bottom panels indicate that plants can switch technologies, which we take into account as well.

We assume that the choice of technology is made at the firm level. The plant manager has to fulfill a production requirement handed down from headquarters. Similarly, we assume that the technology used in the plant is decided before the production year starts, at the firm level. It is the outcome of a net present value comparison between the two technologies. The firm will choose

the technology that gives it the highest discounted profits, taking expectations of future exogenous variables and switching costs into account. Building an empirical model that explicitly makes this net present value comparison is too complicated. A myriad of effects enter this decision. Strategic considerations, pricing of a durable good, and the joint decision for many plants readily come to mind. We would also have to make specific assumptions about the expectations of all exogenous variables and the optimal response functions for all endogenous variables.

Instead, we take a reduced-form approach to control for the technology decision and assume the probability a firm adopts the new technology is a function of observable variables. Crucially, we assume that the technology decision is exogenous to all error terms introduced earlier. This is justifiable, because the technology choice is taken before the start of production. The probability that a plant will enter the sample with the older, traditional technology is given by

$$(1') \quad \psi_{jt} = \frac{1}{1 + \exp(W_{jt}\eta)}.$$

Variables in W capture the relative profitability of each technology. The evolution of the two technologies, experience working with them, and the market segment a plant produces for, are likely to be important determinants. For example, it is argued that lean production is better suited to produce for volatile markets or at a lower scale. The trend towards more cars per-household has led to a larger demand for speciality or niche vehicles, favoring lean technology.

A firm makes the first technology choice before we observe the plant in our sample. In addition, at the start of each year, it has the option of switching technologies. The data spans 27 years allowing 2^{27} possible paths for technology. It is possible to estimate a model that allows plants to shift freely between technologies using the EM algorithm. This necessitates fixing the transition probabilities between technologies to a constant. We opted to be less flexible on the direction of change and more flexible on the transition probabilities, making them vary over time and across plants. We model lean production as an absorbing state, which reduces the number of possible paths for technology to 28. Plants built before the lean technology was available (or actually considered by the firm) can end up with the “wrong” technology for their characteristics. Only these plants will opt to switch and incur switching costs. Using this assumption, we can model the transition probability as a parametric function of observable variables. It also necessitates constructing the likelihood for the entire path of the endogenous variables, instead of year by year.

The transition probabilities are illustrated in the matrix in Table 2. Once a plant adopts the lean technology it is guaranteed to continue with it. Variables in Z determine the probability

Table 2: Transition probabilities for technology

		time $t + 1$	
		T	L
(1)	time t T	$p_{jt} = \frac{1}{1 + \exp(Z_{jt}\gamma)}$	$1 - p_{jt} = \frac{\exp(Z_{jt}\gamma)}{1 + \exp(Z_{jt}\gamma)}$
	L	0	1

of a plant finding it more beneficial to produce with the lean technology rather than sticking with

the traditional one. The same demand variables as in W will influence this transition. In addition, we use a dummy for changeover years to capture switching costs. When a substantially modified model is introduced, a plant has to adjust a large part of the assembly line. This is arguably a good time to make the technology switch as well, because much of the capital stock has to be replaced anyway.

5 Estimation

5.1 Likelihood function

Equations (4b), (4a) and (3) describe the distribution of the endogenous variables —material, output, and the labor-shift ratio— conditional on technology. The probability a plant produces with each technology was derived separately. The vector $y = [\hat{m}, q, \frac{L}{S}]'$ contains the endogenous variables of the system. It depends on two vectors of disturbances $\epsilon_i = [\epsilon_i^m, \epsilon_i^q, \epsilon_i^k]'$, the parameter vectors β_i ($i = T, L$), and the matrix X of all exogenous variables.

We assume that the three errors are independent and normally distributed. Given the nonlinear relation between the labor-shift ration and ϵ_i^k , we need the Jacobian for the transformation of variables,

$$\frac{\partial \epsilon^k}{\partial L/S} = (\frac{L}{S})^{-1}(\lambda_i + \lambda_i^2 - \beta_{il}). \quad (13)$$

The conditions for convexity of the production function also guarantee that the relationship between ϵ_i^k and $\frac{L}{S}$ is monotone. λ_i represents the share of labor in the variable cost, derived earlier. It is a function of observed variables and technology-specific parameters. The density of the y_{jt} -vector conditional on technology i becomes

$$\begin{aligned} h_i(y_{jt}) &= \frac{1}{s_{im}} \phi\left(\frac{\hat{m}_{jt} - g_m(X_{jt}, \beta_i)}{s_{im}}\right) \times \frac{1}{s_{iq}} \phi\left(\frac{q_{jt} - g_q(X_{jt}, \beta_i)}{s_{iq}}\right) \\ &\quad \times \left| \left(\frac{L_{jt}}{S_{jt}}\right)^{-1}(\lambda_{ijt} + \lambda_{ijt}^2 - \beta_{il}) \right| \frac{1}{s_{ik}} \phi\left(\frac{g_i(\frac{L_{jt}}{S_{jt}}) - g_l(X_{jt}, \beta_i)}{s_{ik}}\right), \quad i = T, L \end{aligned}$$

where $\phi(\cdot)$ is the standard normal density. The functional forms for the functions g_m , g_q , g_l , g_T , and g_L are given by equations (4b), (4a), and (3). We make the further assumption that, conditional on the technology choice, errors are uncorrelated over time. The density of y_j for the entire sample becomes¹⁴

$$H(y_{j1}^i \dots y_{jT}^i) = h_i(y_{j1}) \dots h_i(y_{jT}). \quad (14)$$

These densities cannot be estimated directly, if we do not observe the technology choice plants make. For each sequence $(y_{j1}, y_{j2}, \dots, y_{jT})$ we sum over all possible technology paths this can represent. For example, the probability of the sequence $(y_{j1}^T, y_{j2}^T, y_{j3}^L, \dots, y_{jT}^L)$ occurring is $\psi_{j1} p_{j1} (1 - p_{j2})$. The plant started with the traditional production technology, stuck with it after the first year, and switched to lean after the second year. The probability for a sequence y_j —for a plant present in the sample at time 1 and exiting at time T — is

¹⁴We write y_{jt}^i to denote the distribution of y_{jt} conditional on producing with technology i in year t .

$$\begin{aligned}
\mathcal{L}(y_{j1}, y_{j2}, \dots, y_{jT}) = & \quad (1 - \psi_{j1}) \quad H(y_{j1}^L, y_{j2}^L, \dots, y_{jT}^L) \\
& + \psi_{j1}(1 - p_{j1}) \quad H(y_{j1}^T, y_{j2}^L, \dots, y_{jT}^L) \\
& + \psi_{j1}p_{j1}(1 - p_{j2}) \quad H(y_{j1}^T, y_{j2}^T, y_{j3}^L, \dots, y_{jT}^L) \\
& + \dots \\
& + \psi_{j1}p_{j1} \dots (1 - p_{jT-1}) \quad H(y_{j1}^T, y_{j2}^T, \dots, y_{jT-1}^T, y_{jT}^L) \\
& + \psi_{j1}p_{j1} \dots p_{jT-1} \quad H(y_{j1}^T, y_{j2}^T, \dots, y_{jT-1}^T, y_{jT}^T).
\end{aligned} \tag{15}$$

ψ_{j1} is the probability a plant enters the sample with the traditional technology and p_{jt} is the probability a plant remains with the traditional technology in period $t + 1$, given that it had this technology in period t . The starting probability is a function of variables W and the transition probability depends on variables Z . Substituting the previous expressions for $H(\cdot)$ and $h_i(\cdot)$ and multiplying over all plants generates the full likelihood function.

5.2 Results

We estimate equations (4b), (4a), and (3) using the likelihood function in (15). This controls for heterogeneity in technology by jointly estimating the starting and switching probability with the parameters in both production functions. Endogeneity of productivity is accounted for by including plant-fixed effects (not reported) and substituting ϵ_i^k in the production function using the labor-shift equation. All coefficients are estimated jointly and we impose cross-equation restrictions. Results are in Table 3.

The set of coefficients in the first column represents the traditional technology, with standard errors in the second column. The set of coefficients in the third column represents the difference between the parameters for the two technologies. The differences are estimated directly and their standard errors are reported in the fourth column. The coefficients for the lean —absorbing— technology are at the far right, in the fifth column. They are not estimated directly, but obtained by summing the coefficients in the first and third column.

The top panel contains the linear and quadratic input coefficients, the parameter for the returns to shifts, Hicks-neutral and capital-biased productivity growth, and the estimated standard deviation for the shock to production. Most of the standard errors of the difference terms (fourth column) are surprisingly small. This indicates that the two technologies are estimated to be significantly different. The shape of the estimated production functions is interpreted in the next section. We also compare the technologies with the results obtained without controlling for heterogeneity in technology.

Both productivity parameters and the differences are estimated rather precisely and significantly different from zero. The traditional technology is associated with a higher rate of Hicks-neutral productivity growth. The lean technology, on the other hand, experiences a very high rate of capital-biased productivity growth, but Hicks-neutral productivity growth is slightly negative. In section 6 we discuss the importance of these estimates for the industry-wide productivity growth.

The second panel contains the coefficients in the material equation. Although there are no cross-equation restrictions, we estimate it jointly with the other equations, because the technology

Table 3: Estimation results for the three equations and technology choice.

		traditional technology		difference		lean technology
		$\hat{\beta}_T$	(s.e.)	$\hat{\beta}_L - \hat{\beta}_T$	(s.e.)	$\hat{\beta}_L$
Production function:						
shifts	α_{is}	0.9684	(.004)	0.1334	(.034)	1.1017
labor	α_{il}	0.8222	(.023)	0.0224	(.023)	0.8447
capital	α_{ik}	0.1364	(.039)	-0.0325	(.021)	0.1039
labor squared	β_{il}	-0.0148	(.037)	0.0931	(.037)	0.0783
capital squared	β_{ik}	0.0268	(.012)	-0.0204	(.012)	0.0065
labor \times capital	β_{ilk}	0.1483	(.019)	-0.1262	(.019)	0.0221
capital-biased PG	θ_{ik}	0.0706	(.019)	0.2533	(.036)	0.3239
Hicks-neutral PG	θ_{in}	0.0176	(.005)	-0.0269	(.007)	-0.0093
standard deviation	s_{iq}	0.2353	(.009)	0.3132	(.097)	0.5485
Materials equation:						
constant term	α_{i0}	3.6582	(.270)	0.0003	(.313)	3.6586
output	$\frac{1}{\alpha_{im}}$	0.8571	(.022)	0.0099	(.026)	0.8670
quality upgrading	θ_{im}	0.0210	(.002)	-0.0075	(.003)	0.0135
standard deviation	s_{im}	0.2231	(.008)	0.2257	(.008)	0.4488
Labor-shifts equation:						
capital cost	ρ_i	2.9E-6	(5.5E-6)	-2.9E-6	(3.5E-4)	1.4E-11
fixed cost per-shift	μ_i	67.483	(11.4)	-10.021	(11.3)	57.462
union dummy	δ	0.9745	(.123)	X	X	0.9745
standard deviation	s_{iq}	0.0644	(.008)	-0.0340	(.005)	0.0304
Technology Choice:						
Starting probability: η						
constant term		-2.0345	(3.317)			
dummy for cars		0.6853	(2.692)			
dummy for trucks		1.1867	(2.777)			
Japanese dummy		1.0404	(3.742)			
time		0.0911	(0.079)			
Transition probability: γ						
constant term		-1.9412	(0.755)			
dummy for cars		-0.1328	(0.651)			
dummy for trucks		-0.7739	(1.070)			
changeover dummy		0.4080	(1.346)			
time		-0.0383	(0.057)			
number of observations		1358				
likelihood		1890.8				

choice is the same. Quality upgrading through improved components is estimated to be slightly lower for the lean technology. This can be an indication that most quality upgrading was realized before the lean technology became important in the sample. The parameter on output is indistinguishable between technologies. It indicates that an increase in output is accompanied by a less than proportional increase in material input.

The cost-parameters and standard deviation for the labor-shift equation are in the third panel. The parameters determining the marginal product of labor are reported earlier, in the first panel. The dummy for labor unions is assumed constant across technologies. It is estimated close to one, which indicates that unionized plants face almost the same tradeoff between labor and shifts as non-unionized plants. Both the fixed and variable costs related to shifts are estimated to be lower for the lean technology. The variable component is not significantly different from zero for lean producers.

In the bottom panel are the coefficients governing the technology choice. Looking at the starting probability first, the results indicate that plants gradually become more likely to start out with the lean technology. The coefficient on time is not estimated significantly different from zero, but the sign is positive, as expected. Plants that produce only trucks, and to a lesser degree plants producing only cars, prefer the lean technology, although the coefficients are not estimated very precisely. Surprisingly, the new technology seems to be less favored by plants that produce both trucks and automobiles. The average starting probability for the conventional technology is 0.81 at the beginning of the sample period and 0.23 at the end.

Figure 5: Probability that a new plant starts out with the traditional technology.





Figure 5 plots the starting probability for different types of plants. The three dummy variables in Z —only-cars, only-trucks, Japanese ownership— define six types of plants, all of which face a different probability for the traditional technology in each year. Figure 5 only includes the schedule for four of the six types to make the graph better readable. The positive coefficient on time makes the probability for the traditional technology decline over the years. The positive

coefficient estimate on Japanese ownership shifts the schedule for each product down if the plant is owned by a Japanese firm, compared to domestically owned plants. The first Japanese plants entered the sample in 1982.¹⁵

Turning to the transition probability, the negative coefficient on the time variable indicates that the probability for a traditional plant switching to the lean technology declines over the sample period. During a year in which a plant has a major changeover, it is more likely to adopt the new technology. We found earlier that plants producing both cars and trucks had the highest probability to start out with the traditional technology. They also have the highest probability to make the transition to lean technology. Figure 6 traces the inverse of the transition probability, the probability that a traditional plant remains with the traditional technology, in regular and changeover years for the different types of plants. The average probability over the entire sample period is 0.93. This ranges from a high of 0.98 for a truck producer if 1996 was a regular year, to a low of .82 for a car and truck producer if 1963 was a changeover year. These statistics can be interpreted as complements of a hazard rate. The average transition probability suggests that over a 10 year period half of the mass producers would switch to the lean technology. In 1963 it would only take three (and a half) changeover years to achieve the same percentage of mixed-plants switching. A small decrease in the non-transition probability has large effects.

Figure 6: Probability a traditional plants does not switch to the lean technology.



5.3 Interpretation

We draw two conclusions about technology from the estimation results. First, the two technologies are estimated to be rather distinct. The results are consistent with an interpretation of the traditional technology as mass production and the more recent technology as lean production. Secondly, the proportion of plants producing with the traditional technology declines significantly over the

¹⁵The three joint ventures between American and Japanese producers, NUMMI, AutoAlliance, and Diamond-Star, are categorized as Japanese plants

sample period. This is caused by a combination of entry by new plants, which are more likely to be lean, and technology switching by existing plants.

We do find evidence for the existence of two distinct technologies in our sample. The nature of their difference corresponds largely to the mass-lean distinction many industry observers have made. Interpretation of the parameters in the production function directly is complicated by the quadratic terms. We evaluate several statistics at the sample mean, which is calculated separately for each technology. Plants are weighted by the imputed probability for each technology (see below). Table 4 contains estimates for factor shares, returns to scale, productivity growth, and elasticities of substitution. The first two columns contain the results for the two estimated technologies, traditional and lean. For comparison, the far right column presents the same statistics for a translog production function, estimated on the full sample. The translog estimation uses the same output and input measures and includes plant-fixed effects as well.¹⁶

Table 4: Comparing the two technologies with simple translog results

	traditional technology	lean technology	translog estimation for the entire sample
Average factor shares:			
capital share	0.144	0.104	0.349
labor share	0.854	0.764	0.796
labor share in variable cost	0.882	0.695	X
Returns to scale (L & K)	0.998	0.868	1.145
Returns to shifts	0.968	1.102	X
Productivity growth:			
Hicks-neutral	0.018	-0.009	0.004
capital-biased	0.071	0.324	0.007
labor-biased	X	X	-0.006
Elasticity of substitution (L-K)	1.038	1.156	0.334
Demand elasticities:			
ϵ_{LL}	-0.154	-0.884	-0.117
ϵ_{KK}	-0.861	-0.977	-0.266
ϵ_{LK}	0.149	0.120	0.117
ϵ_{KL}	0.886	0.883	0.265

The first two rows contain the labor and capital share evaluated at the mean for each technology. The labor share in variable cost has an additional correction for the returns to shifts. The correct formula and interpretation is given by equation (10). The traditional technology puts more weight on both inputs and produces with constant returns to scale. Returns to shifts are estimated to be constant as well. The lean technology is estimated to have decreasing returns to scale for capital and labor input, but increasing returns with respect to shifts. The lower variable cost of operating a shift (ρ_L was estimated very low) and the increasing returns to shifts both cause lean producers to run more shifts, *ceteris paribus*. Surprisingly, in Figure 3 we found the lean technology to display a higher capital-labor ratio than traditional producers. The relative share of

¹⁶Failure to correct for capacity utilization lead invariably to a negative capital share. We approximate the capital services a plant derives from the observed capital stock, by multiplying the stock with an index of capacity utilization.

capital and labor are estimated rather similar for both technologies, but the higher returns to shifts would lead lean producers to run more shifts and have a lower capital stock per-hour worked (hours worked by employees on different shifts can utilize the same capital stock). This phenomenon can be partly explained by the use of book values for capital. Lean plants are on average more recent and the capital stock has depreciated less.

Hicks-neutral and factor biased productivity estimates differ considerably between the two technologies. Labor-biased productivity growth is restricted to zero, but capital-biased productivity growth is estimated directly. The high rate of capital-biased productivity growth for lean producers is remarkable. Lean production is often associated with flexible machinery. The flexible equipment allows many Japanese plants to produce rather distinct models on the same assembly line. It is also reflected in the lower changeover times for the assembly line between models for Japanese producers. As predicted in Figure 4, the Hicks-neutral productivity growth for the usual translog estimation, with all plants pooled, is estimated between productivity growth for lean and traditional producers. For the translog estimation the factor-biased productivity growth rates do not have the same structural interpretation. The interaction between inputs will affect the productivity estimates as well. In all three columns the same conclusion surfaces: productivity growth is labor-saving.


The bottom panel of Table 4 contains estimates for the elasticity of substitution between capital and labor. The lean technology displays a slightly higher elasticity of substitution and the statistic exceeds one for both technologies. Comparing the results with the translog results indicates that failure to model capacity utilization explicitly leads to a very low elasticity estimate. The same finding is true for the factor-demand elasticities. The lean technology is estimated to be more flexible than the traditional technology. Labor demand, especially, is estimated to be significantly more elastic. Both technologies are estimated to be more price responsive than the results for the simple translog estimation.

The second conclusion we draw from the analysis is that the industry as a whole has almost completed the transition from traditional to lean production. We update the starting probability for the traditional technology for the 78 plants in the sample using the relevant transition probability for each plant. This gives an estimate for the probability a plant is producing with the traditional technology in any given year. Figure 7 shows the evolution of this probability over time. The proportion of the sample producing with the traditional technology clearly declines, but in any given year there is considerable variation. While the probability for the lean technology was very small for all plants in 1963, the reverse is true in 1996.

This finding is caused by two trends. The starting probabilities in Figure 5 already indicated that new plants became much more likely to use the lean technology. At the end of the sample period, every plant built has a lower than 40% probability for the traditional technology, with the probability significantly lower for some types of plants. Out of 49 plants in the sample in 1963 only half remain in 1996. The average entry year for the other half of the plants present at the end of the sample is 1983. In that year the average starting probability for the traditional technology was already below 42%.

The second trend leading to the disappearance of the traditional technology is technology switching by existing plants. For example, car-plants faced a switching probability of around 0.10 for the first part of the sample, even in regular production years. This translates into half of the plants making the technology switch in less than seven years. In addition, switching is more

Figure 7: Probability for traditional technology in the sample



prevalent in changeover years. These changeover years are more common earlier in the sample. The average year for all observations in the data set is 1983, while the average year for a significant model changeover is 1980. At the same time, it is intuitive that the non-transition probability tends to one towards the end of the sample. Gradually less plants are built with the “wrong” technology, necessitating a costly technology switch. This makes the probability for traditional technology for any given plant asymptote to a value above zero, but Figure 7 indicates that this value is rather low for most plants.

Figure 3 provided additional evidence for the substantial change in sample composition. Without identifying individual plants as lean or traditional, the relative weight on the two modes of the distribution shifted considerably from the beginning of the sample period to the end. In particular, the bottom panels isolate the effect of switching, by focussing only on plants that remained in the sample throughout. Both the changing distribution for the capital-labor ratio and the estimated transition probabilities indicate that switching is important.

6 Conclusions about productivity

Finally, we can investigate what we can learn from this about productivity growth in the industry. The estimated productivity growth rates are actual shifts of the production functions, with the function for each technology shifting independently. We find that the traditional technology experiences a higher rate of Hicks-neutral productivity growth, 1.8% versus -0.9% for the lean technology. The lean technology, on the other hand, has a higher capital-biased productivity growth rate, 32.4% versus 7.1%. The capital-biased results are not directly comparable to the usual estimates obtained from a translog cost or production function. The growth rates we estimate only affect total output or costs in proportion to the share of capital and through the interaction between capital and labor. The interpretation is the same though: productivity growth is labor-saving. Because the real

price of labor has a strong upward trend over the sample period, we could expect this finding. It is intuitive to find that the lean technology, which became more popular over the sample period, experiences the higher rate of labor-saving growth.

We can use the estimated technologies to analyze the evolution of productivity in the industry. Figure 1 suggests a break in the trend growth rate for labor productivity in the early 1980s, with productivity growth accelerating strongly. Rather than viewing this as an exogenous shift, we calculate the impact of trends in fundamentals on the industry-wide productivity growth we observe. The first decomposition divides average labor productivity growth in the industry into contributions of lean producers, traditional producers, and a switching effect. Relocation of resources between plants and entry and exit play an important role as well. The second decomposition divides the growth in labor productivity for each technology into the contribution of several effects we estimated, such as total factor productivity growth, returns to scale, and others.

Following Baily et al. (1992) we decompose labor productivity growth for the industry as follows

$$\begin{aligned}\overline{LPG}_t &= \sum_n^{j,k,l} (\theta_{nt} LP_{nt} - \theta_{nt-1} LP_{nt-1}) \\ &= \sum_j^{\text{stay}} \theta_{jt-1} (LP_{jt} - LP_{jt-1}) + \sum_j^{\text{stay}} (\theta_{jt} - \theta_{jt-1}) LP_{jt} + \sum_k^{\text{enter}} \theta_{kt} LP_{kt} - \sum_l^{\text{exit}} \theta_{lt-1} LP_{lt-1}\end{aligned}$$

We use $LP = \log(\frac{Q}{L})$ for the logarithm of the level of labor productivity. The entire expression is the average labor productivity growth, with each plant weighted by its share of the capital stock in the industry. We chose these weights to investigate how capital is reallocated over time.¹⁷ The first term measures the contribution of labor productivity growth at the plant-level, only calculated for plants that stay in the sample from year $t-1$ to year t . The second term measures the relocation effect, also for plants that stayed in the sample. If capital is relocated to plants with above average labor productivity, this term will be positive. The third and fourth term measure the contribution of plants entering and exiting the sample. If new plants are on average more productive than exiting plants, the sum of the last two terms will be positive.

For each plant we decompose the level and growth of labor productivity into the contribution of the traditional and lean technology:

$$LPG_{jt} = \underbrace{\varphi_{jt} LPG_{jt}}_{\text{traditional}} + \underbrace{(1 - \varphi_{jt-1}) LPG_{jt}}_{\text{lean}} + \underbrace{(\varphi_{jt-1} - \varphi_{jt}) LPG_{jt}}_{\text{switching}} \quad (16)$$

$$LP_{nt} = \underbrace{\varphi_{nt} LP_{nt}}_{\text{traditional}} + \underbrace{(1 - \varphi_{nt}) LP_{nt}}_{\text{lean}}, \quad n = j, k, l \quad (17)$$

where φ_{jt} is the probability for the traditional technology for plant j in year t . The first term in (16) captures the contribution to the traditional technology. It multiplies the labor productivity growth with the probability of being traditional in year t . Given that the lean technology is an absorbing state, these plants were also traditional in year $t-1$. The second term measures labor productivity growth times the probability the plant was lean in both years. The third term multiplies the productivity growth with the probability a plant made the technology switch at the start of year

¹⁷For the decomposition, it has to hold that $\sum_j \theta_{jt} + \sum_k \theta_{kt} + \sum_l \theta_{lt} = 1$.

t . Equation (17) similarly decomposes the level of labor productivity into a traditional and lean component.

Substituting equations (16) and (17) in the first decomposition gives nine terms. We sum the contribution of entry and exit for each technology to calculate a net entry effect. Table 5 contains the results for all seven terms, averaged over all years. Because plants differ in their

Table 5: Decomposition of industry-wide labor productivity growth (average 1963-96)

	lean plants	traditional plants	switching
contribution to LPG	1.14%	0.20%	-0.02%
composition	0.06%	-1.02%	
net entry	0.91%	-0.49%	
LPG (full sample)	0.79%		

labor productivity growth, labor productivity level, and the probability for each technology, we find significantly different contributions for both technologies. The lean technology generates most of the industry-wide productivity growth. Plants using this technology increase productivity faster, and new plants are more likely to be lean and more productive than exiting plants. The relocation of capital between plants that stayed in the sample generates a slightly positive contribution for the lean technology, while the mass technology plants relocate inputs to plants with below average productivity. The direct effect of switching is minor. The change in probability for each technology in any given year is small for most plants. In addition, lack of experience is likely to make the newly configured plant operate with lower productivity for some time immediately after a switch in technology.

The second decomposition divides the growth in labor productivity for each technology into the contribution of several fundamental effects that we estimated. We approximate the production function using a Cobb-Douglas function to aggregate labor and capital,

$$Q = S^{\alpha_s} \left[\left(\frac{L}{S} \right)^{\alpha_l} \left(K e^{\theta_k t} \right)^{\alpha_k} \right] e^{\theta_n t}.$$

Taking first differences of the logarithm of this function, deducting growth in labor input, and substituting the coefficient estimates, we write labor productivity growth for each plant as

$$\begin{aligned} LPG = (\dot{q} - \dot{l}) = & \underbrace{\hat{\theta}_n}_{\text{TFPG}} + \underbrace{\hat{\theta}_k \hat{\alpha}_k}_{\text{capital-biased PG}} + \underbrace{\hat{\alpha}_k(\dot{k} - \dot{l} + \dot{s})}_{\text{capital-labor growth}} \\ & + \underbrace{(\hat{\alpha}_l + \hat{\alpha}_k - 1)(\dot{l} - \dot{s})}_{\text{returns to scale}} + \underbrace{(\hat{\alpha}_s - 1)\dot{s}}_{\text{returns to shifts}} + \epsilon, \end{aligned} \quad (18)$$

where \dot{x} denotes the year on year growth rate for x . The last term captures the change in errors between years and the approximation error caused by the use of the Cobb-Douglas approximation.

We sum each of the five terms over all plants, using the probability for each technology as weight. All plant-years receive equal weight, as they did in the estimation. There is no guarantee that the weighted errors sum to zero. To make the decomposition add up, we use the relative importance of the Hicks-neutral and capital-biased productivity growth estimates to divide the

part of labor productivity growth not accounted for by the other three terms. Results for both technologies are in Table 6.

Table 6: Decomposition of labor productivity growth into fundamentals for each technology

	lean technology		traditional technology	
	growth	contribution	growth	contribution
$\overline{\text{LPG}}$	1.46%	(100%)	1.11%	(100%)
TFPG	-0.29%	(-20%)	0.69%	(62%)
capital-biased PG	1.04%	(71%)	0.39%	(36%)
growth in $\frac{K}{L/S}$	0.58%	(39%)	-0.01%	(-1%)
returns to scale	0.08%	(6%)	0.05%	(4%)
returns to shifts	0.05%	(4%)	-0.01%	(-1%)

For the traditional technology, almost all of the growth in labor productivity is accounted for by the estimated total factor and capital-biased productivity growth. Returns to scale are decreasing, but it generates a marginally positive contribution, because the average scale of operation decreased. For the lean technology most of the productivity growth comes through capital-biased productivity growth, unsurprisingly, given the high estimate for this coefficient. In contrast with the traditional technology, there is also a sizable contribution from the increase in capital-labor ratio (for labor per-shift) as well. The returns to scale effect has the same origin as for the traditional technology, but is even more surprising given the sizable diseconomies of scale. Many lean plants have become smaller over time. This is mostly driven by the proliferation of models. This trend in demand added to the popularity of the lean technology, since it necessitates smaller production runs. At the same time, some plants now produce more than one model on the same assembly line, which again favors the more flexible, lean technology.

7 Final Conclusions

We estimated productivity growth using a structural model of production to account for unobservable heterogeneity in technology and productivity. The estimated technologies are consistent with the often made distinction between lean and mass production in this industry. The more recent (lean) technology is associated with higher capital-biased, and lower Hicks-neutral, productivity growth. We also find that the mass production technology is disappearing from the industry. This is caused by the entry of new plants, predominantly choosing the new technology, and technology switching by existing plants.

Using the estimation results, we investigate the trends underlying the large increase in aggregate labor productivity growth for the industry, since the early 1980s. Plant-level growth, attributed to lean producers, and the net entry of plants with the new technology, are the two most important trends. The plant-level labor productivity growth can be further decomposed. We find capital-biased productivity growth and an increase in the capital-labor ratio particularly important. Surprisingly, diseconomies of scale had a positive contribution as well, because plants decreased their scale of operation over the sample period.

Bibliography

Baily Martin N., Charles Hulten, and David Campbell (1992), "Productivity Dynamics in Manufacturing Plants," *Brookings Papers on Economic Activity: Microeconomics*, p. 198-249.

Beard, Randolph T., Steven B. Caudill, and Daniel M. Gropper (1991), "Finite Mixture Estimation of Multiproduct Cost Functions," *Review of Economics and Statistics*, Vol. 73, No. 4, November, p. 654-64.

Blundell, Richard, and Stephen Bond (1998), "GMM Estimation with Persistent Panel Data: An Application to Production Functions," Institute for International Studies Working Paper, No. W99/4.

Caves Douglas W., Laurits R. Christensen and Erwin W. Diewert (1982), "The Economic Theory of Index Numbers and the Measurement of Input, Output, and Productivity," *Econometrica*, Vol. 50, No. 6, p. 1393-1414.

Diamond, Peter, Daniel McFadden, and Miguel Rodriguez (1978), "Measurement of the Elasticity of Factor Substitution and Bias of Technical Change," in Fuss, Melvyn, and Daniel McFadden (eds.), *Production Economics: A Dual Approach to Theory and Applications*, Volume 2, p. 125-147.

Friedlaender Ann F, Clifford Winston, and Kung Wang (1983), "Costs, Technology, and Productivity in the U.S. Automobile Industry," *Bell Journal of Economics*, Vol. 14, No. 1, p. 1-20.

Fuss Melvyn A. and Leonard Waverman (1992), *Cost and Productivity in Automobile Production: The Challenge of Japanese Efficiency*, Cambridge: Cambridge University Press.

Griffith, Rachel (1999), "Using the ARD Establishment Level Data to Look at Foreign Ownership and Productivity in the UK," *Economic Journal*, Vol. 109, No. 456, June, p. 416-42.

Harbour and al. (1999), *The Harbour Report: Competitive Assessment of the North American Automotive Industry*, Rochester.

Klette, T. Jakob, and Zvi Griliches (1996), "The Inconsistency of Common Scale Estimators When Output Prices Are Unobserved and Endogenous," *Journal of Applied Econometrics*, Vol. 11, No. 4, p. 343-61.

Krafcik, John F. (1988), "Comparative Analysis of Performance Indicators at World Auto Assembly Plants," M.A. thesis, MIT.

Mairesse, Jacques, and Zvi Griliches (1990), "Heterogeneity in Panel Data: Are There Stable Production Functions", in ed. Champsaur, Paul et al., *Essays in Honor of Edmond Malinvaud, Volume 3*, Cambridge: MIT Press, p. 192-231.

Milgrom, Paul, and John Roberts (1990), "The Economics of Modern Manufacturing: Technology, Strategy, and Organization," *American Economic Review*, Vol. 80, No. 3, p.511-28.

Olley, G. Steven and Ariel Pakes (1996), "The Dynamics of Productivity in the Telecommunications Equipment Industry," *Econometrica*, Vol. 64, No. 6, p. 1263-97.

Rubenstein, James (1992), *The Changing U.S. Auto Industry: A Geographical Analysis*, New York: Routledge.

Womack, James P., Daniel T. Jones, and Daniel Roos (1990), *The Machine that Changed the World*, New York: Rawson Associates.

Appendix

Property 1 *A higher variable cost for capital (ρ) has an ambiguous effect on investment, but it lowers capacity utilization in almost all cases.*

[Need to add]

Property 2 *High elasticity of substitution between capital and labor increases investment*

[Need to add]

Property 3 *Variation in output over sub-periods increases investment*

If the variable inputs are chosen more frequently than investment, the equation for the optimal capital demand becomes

$$\frac{(\sum_{\tau} E(\frac{\partial G^*(K)}{\partial K}) + r)K}{p^e \sum_{\tau} \bar{Q}_{\tau}} = \sum_{\tau} \frac{Q_{\tau}}{\sum_{\tau} Q_{\tau}} E\left(\frac{\partial \log f_i}{\partial \log K}\right),$$

where τ is the index for sub-periods. On the right-hand side, the capital productivity in each period is weighted by the share of output produced in that period. Because the capital stock is fixed for the entire year, labor input will be higher in periods with high output, which increases capital productivity. If output is spread unevenly capital productivity will be weighted more heavily when it is higher. This increases the right-hand side and therefore the optimal capital stock.